

IN THE CLAIMS:

This listing of claims will replace all prior versions, and listings, of claims in the application:

1. (Previously Presented) A method of congestion management within a switch or network of connected switches wherein the or each of the switches has a plurality of ingress ports and a plurality of egress ports, the method comprising:

when congestion is detected at a first ingress or egress port, sending a message to an upstream port connected to the first ingress or egress port indicating that congestion has occurred at a particular port and requesting storage at the upstream port of data packets destined for that port ; and,

in dependence on the amount of data packets destined for the congested port stored at said upstream port, sending from the upstream port to a further upstream port a message informing said further upstream port of the congestion at the first ingress or egress congested port, said further upstream port storing at said further upstream port data packets destined for the first ingress or egress congested port.

2. (Original) A method according to claim 1, comprising at said upstream port, allocating memory for use as a set-aside- queue for data packets destined for the congested port.

3. (Original) A method according to claim 2, comprising:
at said upstream port creating an entry in a memory to indicate that congestion has occurred at the particular port ; and,

checking packets subsequently received at the upstream port against the entry in the memory and, if a packet is directed to the congested port, storing said packet in the corresponding set aside queue.

4. (Previously Presented) A method according to claim 1, the method comprising within the upstream port, allocating one or more set aside queues in dependence on messages received from the first port.

5. (Original) A method according to claim 4, further comprising within the upstream port controlling data flow into and out of the set aside queue in dependence on the congestion.

6. (Original) A method according to claim 5, comprising de-allocating the one of more set aside queues in dependence on one or more criteria.

7. (Original) A method according to claim 6, in which the one or more criteria include the amount of data in the set aside queue.

8. (Original) A method according to claim 2, in which the message requesting establishment of a set aside queue is discarded by the upstream port if the congestion identified in the request is further downstream than the original congestion.

9. (Previously Presented) A method according to claim 1, in which the message indicating that congestion has occurred includes a token to be kept by the upstream port to identify the upstream port as a leaf port within a congestion tree.

10. (Original) A method according to claim 9, comprising storing data about the number of leaves in the congestion tree in each switch in the tree.

11. (Previously Presented) A method according to claim 9, in which when a set aside queue is de-allocated, the leaf token is returned by the upstream switch to the adjacent downstream switch, the method comprising maintaining a record relating to leaf switches that have returned a leaf token.

12. (Previously Presented) A method according to claim 1, comprising when a subsequent packet is received by the upstream port, if it is destined for the congestion, storing it in a set aside queue, and if it is not destined for the congestion, storing it in a cold queue at the upstream port.

13. (Original) A method according to claim 12, comprising when a packet is received at the upstream port that is destined for the congestion, storing a marker in the cold

queue to provide an indication of the order in which the congestion-bound packet was received with respect to packets already in the cold queue which are also destined for the congestion.

14. (Previously Presented) A method according to claim 3, wherein the memory is provided as an associative memory.

15. (Currently Amended) A method of congestion management within a switch or network of connected switches wherein the or each of the switches has a plurality of ingress ports and a plurality of egress ports, the method comprising:

when congestion is detected at a first ingress or egress port, sending a message to an upstream port connected to the first ingress or egress port indicating that congestion has occurred at a particular port and requesting storage at the upstream port of data packets destined for that port ; and,

in dependence on the amount of data packets destined for the congested port stored at said upstream port, sending from the upstream port to a further upstream port a message informing said further upstream port of the congestion at the first ingress or egress congested port, said further upstream port storing at said further upstream port data packets destined for the first ingress or egress congested port

comprising at said upstream port, allocating memory for use as a set-aside- queue for data packets destined for the congested port.

at said upstream port creating an entry in a memory to indicate that congestion has occurred at the particular port ; and,

checking packets subsequently received at the upstream port against the entry in the memory and, if a packet is directed to the congested port, storing said packet in the corresponding set aside queue

wherein the memory is provided as an associative memory according to claim 14,

in which the associative memory is equipped with a binary command vector operable to engage search logic which in one case is for a set aside queue formation request and thereby performs a minimal length matching operation on the contents of the associative memory and in the case of the assignment of a data packet to a pre-existing set aside queue, thereby performs a maximal length matching operation on the contents of the associative memory.

16. (Original) A method according to claim 15, comprising following receipt of a set aside queue establishment message by the upstream switch, the resulting binary vector that represents the path between the current switch network position and the final congested network destination is left aligned to the index of the current switch position and equipped with a mask, the mask being the size of the bit field describing the route to the congested destination prior to storage in an associative memory element.

17. (Original) A method according to claim 16, whereby data applied to a search register of the associative memory is prior to searching the associative memory elements at its current switch position in a network, left aligned to the index of the current switch network position and equipped with a mask for the purposes of comparison with the stored elements of the associative memory.

18. (Original) A method according to claim 17, wherein a pair of additional inverted bits are used to delineate the start and stop positions of the active section of a turnpool thereby to create a sized mask.

19. (Previously Presented) A signalling protocol for managing congestion within a network of switches, the protocol comprising:

a first message for sending from a first port at which congestion is detected to an upstream port connected to the first port, the first message requesting establishment at the upstream port of a set aside queue for storing data packets received by the upstream switch destined for the source of congestion, the message including a token for storage by said upstream port, the protocol operating such that when said congestion clears, the established set aside queue is de-allocated and the corresponding token is passed downstream in the direction of the previously congested port,

in which when a certain amount of data packets are stored within the set aside queue in said upstream port a message containing a token is sent by said upstream port to a further upstream port requesting establishment of a set aside queue at said further upstream port for storage of data packets destined for the first port at which congestion has been detected.

20. (Original) A protocol according to claim 19, comprising an acknowledgement message for sending from the upstream port to the first port to confirm establishment of the requested set aside queue.

21. (Original) A protocol according to claim 20, comprising a flow control message for sending from the first port to the upstream port including data relating to the congestion at the first port.

22. (Previously Presented) A protocol according to claim 19, comprising a notification for sending from the upstream port to the first port informing the first port of de-allocation of the set aside queue when a set aside queue is no longer required.

23. (Previously Presented) A protocol according to claim 19, comprising a message for informing the first port that the upstream port has de-allocated an old set aside queue.

24. (Previously Presented) A protocol according to claim 19,, comprising a message for sending to the upstream port from the first port instructing the upstream port to modulate its rate of packet transmission to a specified downstream set aside queue.

25. (Cancelled).

26. (Previously Presented) A switch for use in a network of switches, the switch comprising:

two or more ingress ports;

two or more egress ports;

a switch fabric for selectively coupling data packets received at one or more of the ingress ports to one or more of the egress ports;

storage for, in response to a request for storage of data packets destined for a downstream congested port, storing selected data packets;

selection means, for selectively routing a received data packet to the storage in dependence on the detected desired destination of the packet; and

request generation means arranged to send a request to a further upstream port to request storage of data packets destined for the downstream congested port at said further upstream port when a threshold amount of data packets destined for the downstream congested port are stored in the storage.

27. (Original) A switch according to claim 26, in which the selection means comprises a content addressable memory.

28. (Previously Presented) A switch according to claim 26, wherein a set aside queue is only formed in response to the request if one or more of a number of criteria are satisfied.

29. (Currently Amended) A switch for use in a network of switches, the switch comprising:

a plurality of ingress ports for receiving data packets;
a plurality of output ports for transmitting data packets; and,
control means for selectively routing data packets received at one or more of the ingress ports to one or more of the egress ports;

wherein at least one of the ingress ports or egress ports comprises storage for storing details of a congestion tree comprising at least three connected ports in which in use, the switch is located

which at least one of the ingress or egress ports comprises an ingress or egress engine configured in use to receive a data packet; determine from the data packet its eventual destination; and, if the data packet is destined for a congested port to store the packet in the set aside queue and if it is destined for an uncongested port to store the packet in a cold queue for transmission to the uncongested port

the switch being controllable, when connected in a network of switches to execute the method of congestion management within a switch or network of connected switches wherein the or each of the switches has a plurality of ingress ports and a plurality of egress ports, the method comprising:

when congestion is detected at a first ingress or egress port, sending a message to an upstream port connected to the first ingress or egress port indicating that congestion has

occurred at a particular port and requesting storage at the upstream port of data packets destined for that port ; and,

in dependence on the amount of data packets destined for the congested port stored at said upstream port, sending from the upstream port to a further upstream port a message informing said further upstream port of the congestion at the first ingress or egress congested port, said further upstream port storing at said further upstream port data packets destined for the first ingress or egress congested port.

30. (Original) A switch according to claim 29, in which at least one of the ingress or egress ports comprises means for generating a set aside queue for storage of received data packets destined for a port in the congestion tree.

31. (Previously Presented) A switch according to claim 29 in which at least one of the ingress or egress ports is configured in use to generate a set aside queue in response to a request received by the ingress or egress port containing information about congestion at a downstream port, the request containing information about a congested route between the switch and the downstream port.

32. (Cancelled).

33. (Currently Amended) A switch according to claim ~~32~~ 29, in which the ingress or egress engine are embodied in a content addressable memory.

34. (Cancelled).

35. (Previously Presented) A network of interconnected switches connected in a topology, the network comprising a plurality of switches wherein at least two of the switches are switches according to claim 26.

36. (Previously Presented) A signalling protocol for managing congestion within a network of switches, the protocol comprising:

a first message for sending from a first port at which congestion is detected to an upstream port connected to the first port, the first message requesting storage of data packets received by said upstream port destined for the congested first port; and,

a second message for sending by the upstream port to a port further upstream when a threshold amount of data packets destined for the congested first port have been received and stored by the said upstream port, said message requesting storage of data packets destined for the congested first port received by said further upstream port.

37. (Original) A protocol according to claim 36, wherein when storage is requested by either a message from the congested port or the message from said upstream port, said upstream port and said further upstream port respectively are controlled to allocate a set aside queue at said upstream port or at said further upstream port respectively for storage of data packets destined for the congested port.

38. (Original) A protocol according to claim 37, wherein when the set-aside-queue at either or both of said upstream port and said further upstream port have become empty said set-aside-queue may be deallocated.

39. (Previously Presented) An endstation for use in a network of interconnected switches, the end station comprising:

an ingress port for receiving data packets from a network to which in use the end station is connected;

an egress port for providing data packets to a network to which in use the end station is connected;

in which the egress port includes means operable in use to receive a message from a downstream port, the message containing data relating to a congested port further downstream than the downstream port and a request to provide storage for data packets destined for the congested port further downstream.

40. (Original) An endstation according to claim 39, comprising a control device operable in use to, in response to the message received from the network, allocate a set-aside queue for storing of data packets destined for the congested port.

41. (Previously Presented) An endstation according to claim 39 configured for use within the signalling protocol of for managing congestion within a network of switches, the protocol comprising:

a first message for sending from a first port at which congestion is detected to an upstream port connected to the first port, the first message requesting establishment at the upstream port of a set aside queue for storing data packets received by the upstream switch destined for the source of congestion, the message including a token for storage by said upstream port.

42. (Currently Amended) A method of congestion management within a switch or network of connected switches wherein the or each of the switches has a plurality of ingress ports and a plurality of egress ports, the method comprising:

when congestion is detected at a first ingress or egress port, sending a message to an upstream port connected to the first ingress or egress port indicating that congestion has occurred at a particular port and requesting storage at the upstream port of data packets destined for that port ; and,

in dependence on the amount of data packets destined for the congested port stored at said upstream port, sending from the upstream port to a further upstream port a message informing said further upstream port of the congestion at the first ingress or egress congested port, said further upstream port storing at said further upstream port data packets destined for the first ingress or egress congested port according to claim 1,

wherein the step of requesting storage at the upstream port of data packets destined for the congested port comprises requesting establishment of a set aside queue for storage of said data packets; and

wherein data packets stored at said further upstream port are stored in a set aside queue for data packets destined for the congested port thereby establishing an original congestion tree; and

when a subsequent request for storage of data packets is received at any of the ports in the original congestion tree in respect of congestion at a port further downstream than the root of the original congestion tree, accepting the request at the port such that data packets destined for said further downstream port are stored at the port at which the request was received thereby extending the congestion tree downstream.

43. (Original) A method according to claim 42, comprising upon receipt of a request for establishment of a set aside queue at any of said ports in said original congestion tree creating an entry in a memory at the said port to indicate that congestion has occurred at a particular port; and checking data packets subsequently received at the said port against the entry in the memory and, if a data packet is directed to the congested port, storing said data packet in the corresponding set aside queue ; and if a data packet is not directed to the or another congested port, storing the data packet in a cold queue for onward transmission.

44. (Original) A method according to claim 43, comprising: establishing a set aside queue in response to every request for establishment of a set aside queue received at the port, the newly established set aside queue existing concurrently with any already existing set aside queues.

45. (Original) A method according to claim 44, comprising: when a request is for establishment of a set aside queue in respect of a port further downstream than the root of the original congestion tree, placing a link in an existing set aside queue to later activate the newly established set aside queue.

46. (Original) A method according to claim 43, comprising: if a request is for establishment of a set aside queue in respect of a port further upstream than the root of the original congestion tree, overwriting the shortest existing set aside queue with a newly established set aside queue ; and placing a link in the cold queue to the newly established set aside queue.

47. (Original) A method according to claim 46, comprising: if a request is for establishment of a set aside queue in respect of a port further downstream than the root of the original congestion tree, overwriting the existing shortest set aside queue with a newly established set aside queue corresponding to the received request ; and placing a link to the newly established set aside queue in the already existing set aside queue that is the longest already existing set aside queue and that is shorter than the newly established set aside queue.

48. (Previously Presented) A method according to claim 2, comprising, at said further upstream port, allocating memory for use as a set-aside-queue for data packets destined for the first ingress or egress congested port.